Decision Physics and the End of Probabilistic AI Drift

Formalising Determinism in Machine Intelligence

Abstract

Modern artificial intelligence inherits a fundamental instability. Every system based on statistical inference exhibits probabilistic drift: the gradual divergence of outputs from their intended truth over time.

Decision Physics formalises a deterministic alternative. It defines the laws under which identical inputs and states always produce identical outputs, removing stochastic variance from computation.

Deterministic AI replaces probabilistic inference with reproducible computation. It transforms machine intelligence from an empirical discipline into a provable science.

The Reproducibility Crisis in AI

Machine learning is founded on stochastic optimisation. Weight initialisation, random sampling, hardware parallelism and floating-point arithmetic create unpredictable variance between runs. Even with fixed random seeds, two identical models can produce slightly different gradients and, consequently, different decisions.

The field has named these phenomena drift. Data drift arises when the input distribution changes. Concept drift occurs when the meaning of labels evolves. Feedback drift emerges when models alter the very environment they learn from. **Probabilistic drift is intrinsic: the variance created by sampling from probability distributions during inference.**

The result is a reproducibility crisis. Across research and industry, identical experiments do not yield identical results. This undermines regulatory adoption, scientific validation and commercial trust.

The Limits of Current "Determinism"

Developers have attempted to constrain drift through partial methods: seed locking, output caching, and rule-based overlays.

Seed Locking

Fixes the starting point but not the internal arithmetic noise of hardware.

Caching

Repeats previous outputs but hides, rather than removes, nondeterminism.

Rule-Based Systems

Guarantee repeatability only by abandoning learning altogether.

None of these approaches meet regulatory-grade determinism. They cannot provide the verifiable audit trails required in finance, healthcare or law. They stabilise symptoms, not the system.

Decision Physics: Formal Laws of Determinism

Decision Physics defines the mathematical conditions that make machine intelligence reproducible. It operates through four interdependent laws:

DP-1 Replay Invariance

Identical inputs and states always yield identical outputs.

DP-2 Symbolic Isomorphism

Semantically equivalent expressions resolve to identical canonical forms.

DP-3 Lineage Conservation

Every output carries an immutable provenance record Λ that documents its computational history.

DP-4 Drift Nullification

Output stability persists regardless of external model updates, corpus changes, or hardware variation.

Together these laws transform probabilistic inference $p(y|x)$ into deterministic computation $y = F(x, state)$.

Under Decision Physics, intelligence becomes a function, not a distribution.

The Mechanism (Abstracted)

Decision Physics is implemented through the CML-H pipeline:

1. **Input**

   Raw data entry

2. **Typed Primitives**

   Encode input data with fixed semantic types

3. **Canonical Form CF(P)**

   Normalise structure and meaning into standard representation

4. **Stable Digest H**

   Apply deterministic hashing to produce reproducible signature

5. **Lineage Λ**

   Store all transformation steps for audit and verification

Symbolic canonicalisation ensures that semantically identical content, even across languages, generates the same digest H in accordance with DP-2. This allows reproducibility not only of data but of meaning.

Implementation Proof: Deterministic Build System

The principles of Decision Physics were tested through our MatrixOS and Decision Physics Gate.

Within this environment, clocks are synchronised to deterministic intervals, random number generators are seeded and locked through verifiable receipts, and all arithmetic is normalised to fixed precision. Each computational event is logged and hashed into its lineage record.

1,000

Identical Runs

Produced 1,000 identical outputs

0

Drift Detected

No drift across time, hardware or state

100%

Verifiable

Each output verified by SHA digests and lineage comparison

This establishes empirical proof that Decision Physics enforces true replay invariance.

Comparative Analysis

Property	Probabilistic AI	Deterministic AI
Output consistency	Variable	**Fixed**
Explainability	Post-hoc	**Intrinsic**
Auditability	Partial	**Complete**
Compliance	Risk-laden	**Defensible**
Scaling risk	Compounds	**Nullifies**

Deterministic AI provides reliability at every level of deployment, from experimental replication to enterprise governance.

Empirical Cases

Finance

Deterministic modelling prevents cascading failures and black-swan amplification. Regulatory bodies can verify every output against its lineage.

Pharmaceutical Research

Trial simulations become reproducible. Each prediction and dataset transformation is documented and repeatable, accelerating approval cycles.

Software as a Service

Conversion engines maintain stable behaviour across updates, eliminating drift-induced revenue variance.

Government and Policy

Automated systems generate decisions that can be traced, reviewed and defended, restoring accountability to algorithmic governance.

Theoretical Implications and Conclusion

Decision Physics collapses the probabilistic ceiling. The structural thirty per cent error margin that defines current AI performance reduces to zero.

Artificial intelligence becomes a computational science rather than an empirical art. Models are no longer trained to approximate truth but to compute it reproducibly.

This foundation enables deterministic regulation, deterministic valuation and deterministic trust infrastructure. It transforms AI from a discipline of prediction into a framework of proof.

> Deterministic AI is not a control layer placed on top of existing systems. It is a new physics for intelligent computation.

Decision Physics formalises the principles that make intelligence reproducible, explainable and stable. It replaces stochastic approximation with certainty.

The proof is operational. MatrixOS's deterministic builds have demonstrated zero drift and full lineage verification.

Decision Physics is testable. The age of probabilistic drift is closing. The era of deterministic intelligence has begun.